

Machine Learning for NLP

Lecture 2: Linear classifiers




**UNIVERSITY OF
GOTHENBURG**

Richard Johansson

September 2, 2016

vectors

- ▶ a tuple consisting of n numbers is called a **vector**
- ▶ the set of all possible tuples of length n is called an n -dimensional **vector space**
- ▶ for instance: $(1, 2)$ is a 2-dimensional vector
- ▶ they can be interpreted geometrically, either as a point in a coordinate system



- ▶ ...or as a direction (e.g. of motion or force)

NumPy linear algebra examples

```
>>> import numpy
>>> v1 = numpy.array([1, 0, 0, 1, 0])
>>> v2 = numpy.array([0, 2, 1, -2, 1])
>>> v1
array([1, 0, 0, 1, 0])
>>> v2
array([ 0,  2,  1, -2,  1])
>>> v1 + v2
array([ 1,  2,  1, -1,  1])
>>> 100 * v1
array([100,   0,   0, 100,   0])
>>> numpy.dot(v1, v2)
-2
>>> v1.dot(v2)
-2
>>> numpy.linalg.norm(v1)
1.4142135623730951
```


matrix multiplication

- ▶ matrix **multiplication** is an extension of the dot product for vectors
- ▶ each cell in the new matrix is computed as the dot product between a row and a column:

$$\begin{bmatrix} 1 & 2 \\ 3 & 4 \end{bmatrix} \cdot \begin{bmatrix} 10 & 20 \\ 30 & 40 \end{bmatrix} = \begin{bmatrix} 70 & 100 \\ 100 & 130 \end{bmatrix}$$


geometric interpretation of matrix multiplication

- ▶ as mentioned, we use matrix multiplication (and other matrix operations) mainly for efficiency in this course
 - ▶ a matrix multiplication instead of many dot products
- ▶ however, in geometry we can use matrix multiplication can be used to express many useful transformations
 - ▶ scaling
 - ▶ rotation
 - ▶ projection from 3D to 2D
 - ▶ ...

the first step: mapping features to numerical vectors


- ▶ scikit-learn's learning methods works with features as **numbers**, not strings
- ▶ they can't directly use the feature dicts we have stored in X
- ▶ converting from string to numbers is the purpose of these lines:

```
vec = DictVectorizer()  
Xe = vec.fit_transform(X)
```




types of vectorizers

- ▶ a DictVectorizer converts from attribute–value dicts:




- ▶ a CountVectorizer converts from texts (after applying a tokenizer) or lists:



- ▶ a TfidfVectorizer is like a CountVectorizer, but also uses TF*IDF


what goes on in a DictVectorizer?

- ▶ each feature corresponds to one or more columns in the output matrix
- ▶ easy case: boolean and numerical features:




what goes on in a DictVectorizer?

- ▶ each feature corresponds to one or more columns in the output matrix
- ▶ easy case: boolean and numerical features:



- ▶ for string features, we reserve one column for each possible value: **one-hot encoding**
 - ▶ that is, we convert to booleans



code example (DictVectorizer)

```
from sklearn.feature_extraction import DictVectorizer
X = [{'f1':'NP', 'f2':'in', 'f3':False, 'f4':7},
     {'f1':'NP', 'f2':'on', 'f3':True, 'f4':2},
     {'f1':'VP', 'f2':'in', 'f3':False, 'f4':9}]
vec = DictVectorizer()
Xe = vec.fit_transform(X)
print(Xe.toarray())

print(vec.vocabulary_)
```

code example (DictVectorizer)

```
from sklearn.feature_extraction import DictVectorizer
X = [{ 'f1': 'NP', 'f2': 'in', 'f3': False, 'f4': 7},
      { 'f1': 'NP', 'f2': 'on', 'f3': True, 'f4': 2},
      { 'f1': 'VP', 'f2': 'in', 'f3': False, 'f4': 9}]
vec = DictVectorizer()
Xe = vec.fit_transform(X)
print(Xe.toarray())

print(vec.vocabulary_)
```


the result:

```
[[ 1.  0.  1.  0.  0.  7.]
 [ 1.  0.  0.  1.  1.  2.]
 [ 0.  1.  1.  0.  0.  9.]]
```

```
{'f4': 5, 'f2=in': 2, 'f1=NP': 0, 'f1=VP': 1, 'f2=on': 3, 'f3': 4}
```


CountVectorizers for document representation

- ▶ a `CountVectorizer` converts from documents
 - ▶ the document is a string or a list of tokens
- ▶ just like string features in a `DictVectorizer`, we use one-hot encoding so that each word type will correspond to one column



code example (CountVectorizer)

```
X = ['example text',  
     'another text']  
  
vec = CountVectorizer()  
Xe = vec.fit_transform(X)  
print(Xe.toarray())  
  
print(vec.vocabulary_)
```

code example (CountVectorizer)


```
X = ['example text',  
     'another text']  
  
vec = CountVectorizer()  
Xe = vec.fit_transform(X)  
print(Xe.toarray())  
  
print(vec.vocabulary_)
```

the result:

```
[[0 1 1]  
 [1 0 1]]  
  
{'text': 2, 'example': 1, 'another': 0}
```




geometric view

- ▶ geometrically, a linear classifier can be interpreted as separating the vector space into two regions with a line (plane, hyperplane)



linear separability


- ▶ a dataset is **linearly separable** if there exists a \mathbf{w} that gives us perfect classification



- ▶ theorem: if the dataset is linearly separable, then the perceptron learning algorithm will find a separating \mathbf{w} in a finite number of steps

a simple example of linear inseparability

<i>very good</i>	Positive
<i>very bad</i>	Negative
<i>not good</i>	Negative
<i>not bad</i>	Positive




mapping into a larger vector space

- ▶ we may add **combinations** of features to make the dataset separable:

<i>very good</i>	<i>very-good</i>	Positive
<i>very bad</i>	<i>very-bad</i>	Negative
<i>not good</i>	<i>not-good</i>	Negative
<i>not bad</i>	<i>not-bad</i>	Positive

- ▶ from a geometrical viewpoint: we are creating a feature space with a higher dimensionality:



- ▶ lots of features → LOTS of combinations

coding a linear classifier using NumPy

```
class LinearClassifier(object):  
    def predict(self, x):  
        score = x.dot(self.w)  
        if score >= 0.0:  
            return self.positive_class  
        else:  
            return self.negative_class
```

better: handle all instances at the same time

```
class LinearClassifier(object):  
    def predict(self, X):  
        scores = X.dot(self.w)  
        out = numpy.select([scores>=0.0, scores<0.0],  
                           [self.positive_class,  
                            self.negative_class])  
  
        return out
```

an illustration of the steps

```
>>> import numpy

>>> scores = numpy.array([-1, 2, 3, -4, 5])

>>> scores >= 0
array([False,  True,  True, False,  True], dtype=bool)

>>> scores < 0
array([ True, False, False,  True, False], dtype=bool)

>>> numpy.select([scores >= 0, scores < 0], ["positive", "negative"])
array(['negative', 'positive', 'positive', 'negative', 'positive'],
      dtype='<S8')
```

perceptron reimplementation in NumPy

```
class NewPerceptron(LinearClassifier):

    def __init__(self, n_iter=10):
        self.n_iter = n_iter

    def fit(self, X, Y):
        # ... some initialization

        X = X.toarray() # convert sparse to dense
        n_features = X.shape[1]
        self.w = numpy.zeros( n_features )

        for i in range(self.n_iter):
            for x, y in zip(X, Y):

                score = self.w.dot(x)

                if score <= 0 and y == self.positive_class:
                    self.w += x
                elif score >= 0 and y == self.negative_class:
                    self.w -= x
```


reimplementation with sparse vectors

```
class SparsePerceptron(LinearClassifier):  
  
    # ...  
  
    def fit(self, X, Y):  
        # ... some initialization  
  
        for i in range(self.n_iter):  
            for x, y in zip(X, Y):  
  
                score = sparse_dense_dot(x, self.w)  
  
                if y*score <= 0:  
                    add_sparse_to_dense(x, self.w, y)
```

comparison

- ▶ on my computer, with the data set we'll use in assignment 2:
 - ▶ dense vectors: 17 seconds
 - ▶ sparse vectors: 3 seconds

